

RESPONSE-ADAPTIVE DOSE-FINDING UNDER MODEL UNCERTAINTY

BY BJÖRN BORNKAMP, FRANK BRETZ, HOLGER DETTE¹
AND JOSÉ PINHEIRO

*TU Dortmund University, Novartis Pharma AG, Ruhr University
at Bochum and Johnson & Johnson*

Dose-finding studies are frequently conducted to evaluate the effect of different doses or concentration levels of a compound on a response of interest. Applications include the investigation of a new medicinal drug, a herbicide or fertilizer, a molecular entity, an environmental toxin, or an industrial chemical. In pharmaceutical drug development, dose-finding studies are of critical importance because of regulatory requirements that marketed doses are safe and provide clinically relevant efficacy. Motivated by a dose-finding study in moderate persistent asthma, we propose response-adaptive designs addressing two major challenges in dose-finding studies: uncertainty about the dose-response models and large variability in parameter estimates. To allocate new cohorts of patients in an ongoing study, we use optimal designs that are robust under model uncertainty. In addition, we use a Bayesian shrinkage approach to stabilize the parameter estimates over the successive interim analyses used in the adaptations. This approach allows us to calculate updated parameter estimates and model probabilities that can then be used to calculate the optimal design for subsequent cohorts. The resulting designs are hence robust with respect to model misspecification and additionally can efficiently adapt to the information accrued in an ongoing study. We focus on adaptive designs for estimating the minimum effective dose, although alternative optimality criteria or mixtures thereof could be used, enabling the design to address multiple objectives. In an extensive simulation study, we investigate the operating characteristics of the proposed methods under a variety of scenarios discussed by the clinical team to design the aforementioned clinical study.

Received April 2010; revised November 2010.

¹Supported in part by the Collaborative Research Center “Statistical modeling of non-linear dynamic processes” (SFB 823) of the German Research Foundation (DFG).

Key words and phrases. Dose-response, drug development, minimum effective dose, optimal design, shrinkage approach.

This is an electronic reprint of the original article published by the [Institute of Mathematical Statistics](#) in *The Annals of Applied Statistics*, 2011, Vol. 5, No. 2B, 1611–1631. This reprint differs from the original in pagination and typographic detail.

1. Introduction. Dose-finding studies have several challenges in common. First, they usually address two distinct objectives, which lead to different requirements on the study design [Ruberg (1995), Bretz et al. (2008)]: (i) assessing evidence of a drug effect, and (ii) estimating relevant target doses. Second, the form of the dose-response relationship is unknown prior to the study, leading to model uncertainty. This problem is often underestimated, although ignoring model uncertainty can lead to highly undesirable effects [Chatfield (1995), Draper (1995), Hjorth (1994)]. Third, data from dose-finding studies are usually highly variable. This issue is of particular importance in pharmaceutical drug development, because sample sizes are kept to a minimum for ethical and financial reasons. It is therefore critical to develop efficient dose-finding study designs that use the limited information as efficiently as possible, while addressing the above challenges.

Many approaches have been proposed in the optimal design literature to distribute patients efficiently with regard to given study objectives; see Wu (1988), Fedorov and Leonov (2001) and King and Wong (2004), among many others. However, most of this work has concentrated on an assumed fixed dose-response model. As there is typically considerable model uncertainty at the planning stage of a dose-response study, these methods have limited practical use. Based on concepts introduced by Läuter (1974) [see also Cook and Wong (1994), Zhu and Wong (2000, 2001), Biedermann, Dette and Pepelyshev (2006)], Dette et al. (2008) investigated model-robust designs that provide efficient target dose estimates for a set of candidate dose-response models, rather than for a single dose-response model. However, their designs require knowledge about the unknown parameters associated with the anticipated dose-response models as well as the prespecification of model probabilities.

A natural remedy is to investigate response-adaptive designs (adaptive designs, in short) with several cohorts of subjects. After each stage the accumulated information of the ongoing study is used to update the initial information of the underlying model parameters and model probabilities, which in turn is used to calculate the design for the subsequent stage(s). Several adaptive designs have been developed for this problem; see, for example, Miller, Guilbaud and Dette (2007) and Dragalin, Hsuan and Padmanabhan (2007) for recent approaches using optimal design theory, or Zhou et al. (2003), Müller et al. (2006), and Wathen and Thall (2008) for recent Bayesian adaptive designs. Dragalin et al. (2010) performed an extensive simulation study that compared five different adaptive dose-finding methods.

In this paper we propose adaptive designs addressing the three major challenges described above: multiple study objectives, model uncertainty and large variability in the data. For this purpose we use the model-robust designs proposed by Dette et al. (2008) together with a Bayesian shrinkage

approach to stabilize the parameter estimates, especially in the early part of a study. This allows one to calculate parameter estimates as well as model probabilities that can then be used to calculate model-robust designs for the subsequent stage(s) of the study. The resulting designs are robust with respect to model misspecification and additionally adapt to the continuously accrued information in an ongoing study. We focus on adaptive designs for estimating the minimum effective dose (MED), that is, the smallest dose achieving a clinically relevant benefit over the placebo response. However, alternative optimality criteria or mixtures of optimality criteria could be used, enabling the design to address multiple objectives.

2. Asthma dose-finding study. The research for this article was motivated by a Phase II dose-finding study for the development of a new pharmaceutical compound in asthma. This was a multi-center, randomized, double-blind, placebo controlled, parallel group study in patients with moderate persistent asthma, who were randomized to one of seven active dose levels or placebo. The primary endpoint was change from baseline in a lung function parameter (forced expiratory volume in 1 second, FEV_1) after 28 days of administration, scaled such that larger values indicated a better outcome. The objective of the trial was to evaluate the dose effects over placebo for the primary endpoint and to assess whether there was any evidence of a drug effect. Once such a dose-response signal had been detected, one would subsequently estimate relevant target doses, where the primary focus was on estimating the MED.

Based on discussions with the clinical team, a homoscedastic normal model was assumed for the primary endpoint with a standard deviation of 350 ml, a placebo effect of 100 ml and a maximum treatment effect of 300 ml within the dose range $[0, 50]$ under investigation. The available doses were 0 (= placebo), 0.5, 1.0, 2.5, 5, 10, 20 and 50. The clinically relevant benefit over the placebo effect was set to 200 ml. That is, an increase in treatment effect of less than 200 ml over the observed placebo response was considered to be clinically irrelevant. Furthermore, all dose levels within the investigated dose range were considered safe based on previous studies, so that efficacy was of primary interest.

Because this study was conducted early in the drug development program, limited information about the dose-response shape was available at the planning stage. A set of candidate dose-response models was derived before starting the study; see Table 1 and Figure 1 for the full model specifications (including a preliminary specification of the model parameters). An increase of the dose-response curve in the lower part of the investigated dose range was considered likely, so two concave increasing models (E_{\max_1} , E_{\max_2}) were included in the model set. In addition, S -shaped (Logistic_1),

unimodal (Beta) and convex (Logistic₂) models were included in the candidate model set to robustify the statistical analysis with respect to model uncertainty. We refer to Pinheiro, Bornkamp and Bretz (2006) for details on the use of candidate models in dose-response studies and the elicitation of best guesses for the model parameters.

TABLE 1
Candidate dose response models as a function of dose d , where
 $B(a, b) = (a + b)^{a+b} / (a^a b^b)$

Model	Full model specification	Model parameters	True MED
Beta	$\theta_0 + \theta_1 B(\theta_2, \theta_3)(d/60)^{\theta_2}(1 - d/60)^{\theta_3}$	(100, 300, 0.43, 0.6)	5.21
E _{max} ₁	$\theta_0 + \theta_1 d / (\theta_2 + d)$	(100, 420, 20)	18.18
E _{max} ₂	$\theta_0 + \theta_1 d / (\theta_2 + d)$	(100, 330, 5)	7.69
Logistic ₁	$\theta_0 + \theta_1 / \{1 + \exp[(\theta_2 - d)/\theta_3]\}$	(98, 302, 17.5, 3.3)	19.82
Logistic ₂	$\theta_0 + \theta_1 / \{1 + \exp[(\theta_2 - d)/\theta_3]\}$	(92, 615, 50, 11.5)	42.28

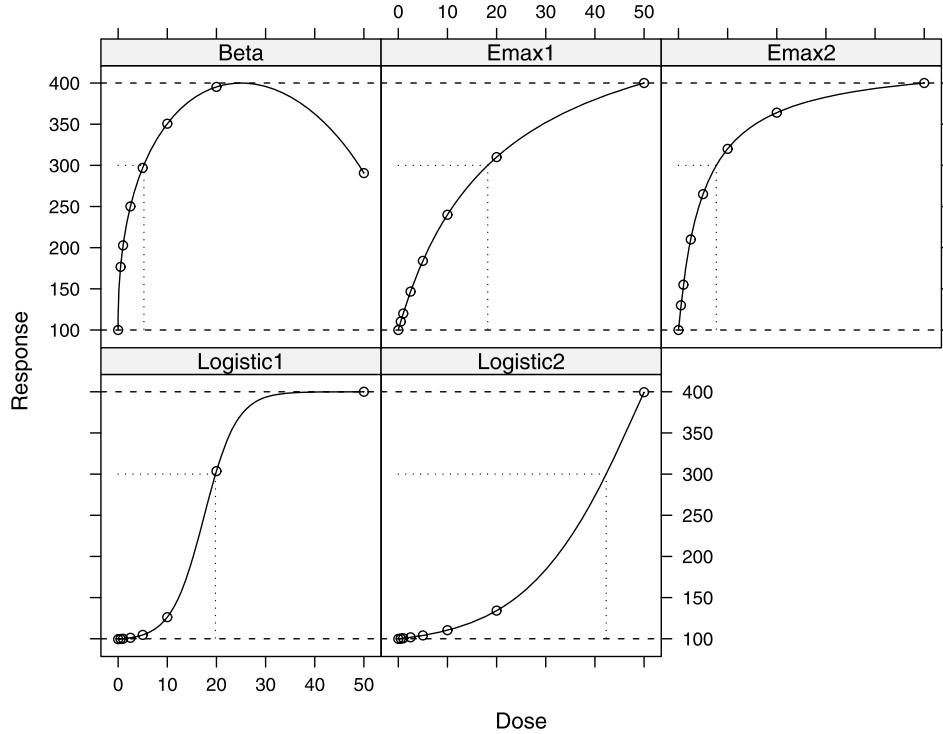


FIG. 1. Graphical display of the dose-response models in Table 1. Open dots denote the potential responses at the seven active dose levels and placebo available in the study. Dotted horizontal lines indicate the clinical relevance threshold on top of placebo response and dotted vertical lines the resulting MED.

Given the information and constraints above, the clinical team was faced at the planning stage with several remaining key questions on the study design:

- (A) Should an adaptive design be employed at all or would a nonadaptive design be sufficient?
- (B) If the decision was to employ an adaptive design, how many interim analyses should be conducted?
- (C) How many dose levels should be included in the study, that is, are all seven active dose levels from above needed?
- (D) If not all active dose levels were needed, which of them should then be investigated?

In addition to these statistical questions, many further considerations were discussed by the clinical team: adaptive designs require more logistical effort to set up the repeated data collection and cleaning/analysis processes than nonadaptive designs; including all seven active doses in the study would pose serious challenges to the drug manufacturing and supply departments, especially if the allocation changed during the study course; and how to ensure trial integrity and validity. In the following we focus on the statistical questions and describe the proposed methodology for the study.

3. Methodology. Assume k distinct dose levels d_1, \dots, d_k , where $d_1 = 0$ denotes placebo. Let n_i patients be allocated to dose d_i and $N = \sum_{i=1}^k n_i$. The vector of allocation weights is denoted by $\mathbf{w} = (w_1, \dots, w_k)'$, where $w_i = n_i/N$. Let further $Y_{ij} \sim N(f(d_i, \boldsymbol{\theta}), \sigma^2)$ denote the observation of patient $j = 1, \dots, n_i$ at dose $d_i, i = 1, \dots, k$, where the dose-response model $f(\cdot)$ is parameterized through the parameter vector $\boldsymbol{\theta}$ and $N(\mu, \sigma^2)$ denotes the normal distribution with mean μ and variance σ^2 .

Most dose-response models used in practice, including those in Table 1, can be decomposed as

$$(3.1) \quad f(d, \boldsymbol{\theta}) = \theta_0 + \theta_1 f^0(d, \boldsymbol{\theta}^0),$$

where $\boldsymbol{\theta} = (\theta_0, \theta_1, \boldsymbol{\theta}^{0'})' = (\theta_0, \dots, \theta_p)'$. The parameters $\boldsymbol{\theta}^* = (\theta_0, \theta_1)'$ enter the model function f linearly and determine its location and scale, while f^0 is typically a nonlinear function that determines the shape of the model function f through the parameters $\boldsymbol{\theta}^0$.

The minimum effective dose producing a clinically relevant effect Δ over the placebo response is defined as

$$(3.2) \quad \text{MED} = \min_{d \in (d_1, d_k]} \{f(d, \boldsymbol{\theta}) > f(d_1, \boldsymbol{\theta}) + \Delta\},$$

where we assume that a beneficial effect is associated with larger values of the response variable. Note that the MED may not exist, as no dose in $(d_1, d_k]$ may produce an improvement of Δ compared with placebo.

3.1. *Robust designs for MED estimation.* Given a function f^0 , it follows from (3.2) that the MED (provided it exists) is a solution to

$$(3.3) \quad \theta_0 + \theta_1 f^0(0, \boldsymbol{\theta}^0) + \Delta = \theta_0 + \theta_1 f^0(\text{MED}, \boldsymbol{\theta}^0).$$

Consequently, $\text{MED} = b(\boldsymbol{\theta}) = h^0(f^0(0, \boldsymbol{\theta}^0) + \Delta/\theta_1)$, where $h^0(x) = \inf\{z | f^0(z) \geq x\}$ denotes the (generalized) inverse of the function f^0 with respect to the variable d . Standard asymptotic theory for nonlinear models [Seber and Wild (1989)] yields that the maximum likelihood (ML) estimate $\hat{\boldsymbol{\theta}}$ is approximately multivariate normal distributed with mean vector $\boldsymbol{\theta}$ and covariance matrix $\frac{\sigma^2}{N} \mathbf{M}^{-1}(\boldsymbol{\theta}, \mathbf{w})$, where $\mathbf{M}(\boldsymbol{\theta}, \mathbf{w}) = \sum_{i=1}^k w_i g(d_i, \boldsymbol{\theta}) g(d_i, \boldsymbol{\theta})'$ denotes the information matrix and $g(d, \boldsymbol{\theta}) = \frac{\partial f(d, \boldsymbol{\theta})}{\partial \boldsymbol{\theta}}$ the gradient of the dose-response model f with respect to $\boldsymbol{\theta}$. It follows from the δ -method [see van der Vaart (1998)] that the MED estimator based on $\hat{\boldsymbol{\theta}}$, $\widehat{\text{MED}} = b(\hat{\boldsymbol{\theta}})$, is asymptotically normally distributed with mean $b(\boldsymbol{\theta})$ and variance $V(\boldsymbol{\theta}, \mathbf{w}) = \frac{\sigma^2}{N} \nabla b(\boldsymbol{\theta})' \mathbf{M}^{-1}(\boldsymbol{\theta}, \mathbf{w}) \nabla b(\boldsymbol{\theta})$, where $\nabla b(\boldsymbol{\theta}) = \frac{\partial b(\boldsymbol{\theta})}{\partial \boldsymbol{\theta}}$. Hence, minimizing $V(\boldsymbol{\theta}, \mathbf{w})$ with respect to $\mathbf{w} \in \mathbb{S}^k = \{\mathbf{w} | \sum_{i=1}^k w_i = 1, \mathbf{w} \geq 0\}$ results in optimal designs that minimize the approximate variance of $\widehat{\text{MED}}$. This design criterion has also an appealing decision theoretic justification: The asymptotic normal distribution of $\widehat{\text{MED}}$ approximates the posterior distribution of the MED in a Bayesian model framework. Hence, minimizing the log-variance of $\widehat{\text{MED}}$ is equivalent to minimizing the (approximate) Shannon entropy of the posterior distribution of the MED [Chaloner and Verdinelli (1995)].

In principle, the above optimization could be done with respect to the number and choice of doses and their corresponding allocation ratios [Dette et al. (2008)], but, in practice, manufacturing constraints often determine the available doses, as it was the case in the asthma study from Section 2. In the following we thus restrict the optimization to the weights \mathbf{w} for prespecified doses d_1, \dots, d_k .

The true dose-response function f is unknown and optimal designs are typically not robust with respect to model misspecification [Dette et al. (2008)]. In the following we assume a set of M candidate models $f_m(d, \boldsymbol{\theta}_m) = \theta_{0m} + \theta_{1m} f_m^0(d, \boldsymbol{\theta}_m^0)$, $m = 1, \dots, M$, such as those described in Table 1. We “integrate” the design criterion conditional on model m with respect to the model probabilities α_m . Hence, using the design criterion $\sum_{m=1}^M \alpha_m \log(V_m(\boldsymbol{\theta}_m, \mathbf{w}))$ or, equivalently,

$$(3.4) \quad \Psi(\mathbf{w}) = \prod_{m=1}^M (V_m(\boldsymbol{\theta}_m, \mathbf{w}))^{\alpha_m}$$

leads to designs that are robust with respect to model misspecification, where $V_m(\boldsymbol{\theta}_m, \mathbf{w})$ denotes the variance of the estimate for the MED in the m th

model ($m = 1, \dots, M$). Note that because of taking logarithms above, there is no need to standardize the individual model variances. Otherwise this would be necessary to avoid that some models dominate the design criterion [the $V_m(\boldsymbol{\theta}_m, \mathbf{w})$ can be quite model dependent and differing in size]. However, the numerical calculation of robust designs using the criterion (3.4) requires the knowledge of $\boldsymbol{\theta}_m$ and $\alpha_m, m = 1, \dots, M$. In the following sections we describe how the initial best parameter guesses can be updated during an ongoing study such that subsequent stages can be redesigned based on the updated estimates for $\boldsymbol{\theta}_m$ and $\alpha_m, m = 1, \dots, M$; see Section 3.3 for a description of the complete procedure in an algorithmic form.

3.2. Updating of model parameters and weights. Reliably estimating the parameters $\boldsymbol{\theta}_1, \dots, \boldsymbol{\theta}_M$ is a challenging problem, particularly in early stages of a study. ML estimates for these parameters are typically highly variable, and may even not exist without imposing bounds on the parameter space. One way of stabilizing estimates is to use a shrinkage approach based on, for example, penalized maximum likelihood or maximum a-posteriori (MAP) estimates. Here, one optimizes the log-likelihood function plus a term which determines the prior plausibility of the parameters (the log prior distribution). The estimate is then a compromise between the information contained in data and the prior distribution. This stabilizes the estimates in early stages due to the shrinkage toward a priori reasonable values. In later stages the shrinkage effect decreases because the log prior remains constant while the log likelihood receives more weight with increasing sample sizes. If a completely flat prior distribution is used, standard ML and MAP estimation coincide, so that using nonuniform priors is desirable. We discuss the choice of nonuniform priors in more detail further below.

Apart from stable parameter estimates $\boldsymbol{\theta}_1, \dots, \boldsymbol{\theta}_M$ for the dose-response models, one needs to update the model probabilities $\alpha_1, \dots, \alpha_M$ at an interim analysis. We propose using a probability distribution over the different dose-response models and evaluating the posterior probabilities for each model after having observed the data; see, for example, Kass and Raftery (1995) for a detailed description of posterior probabilities and Bayes factors. These posterior model probabilities can then be used in the design criterion (3.4). A computationally efficient approach to approximate the posterior model probabilities is the Bayesian information criterion (BIC). However, previous simulation studies in the context of dose-finding studies showed that the BIC approximation frequently favors too simplistic models for realistic variances and sample sizes [see Bornkamp (2006)]. Other approximate methods, such as fractional Bayes factors, or intrinsic Bayes factors [see O’Hagan (1995) or Berger and Pericchi (1996)], either depend on arbitrary tuning parameter values or are computationally prohibitive. Thus, for each model we will use the exact posterior probabilities resulting from the prior distributions

assumed for the MAP estimation. In our case, these probabilities can be calculated using efficient numerical quadrature without the need to resort to computationally expensive Markov chain Monte Carlo techniques. In the remainder of this section we provide details on the prior elicitation and the calculation of posterior probabilities.

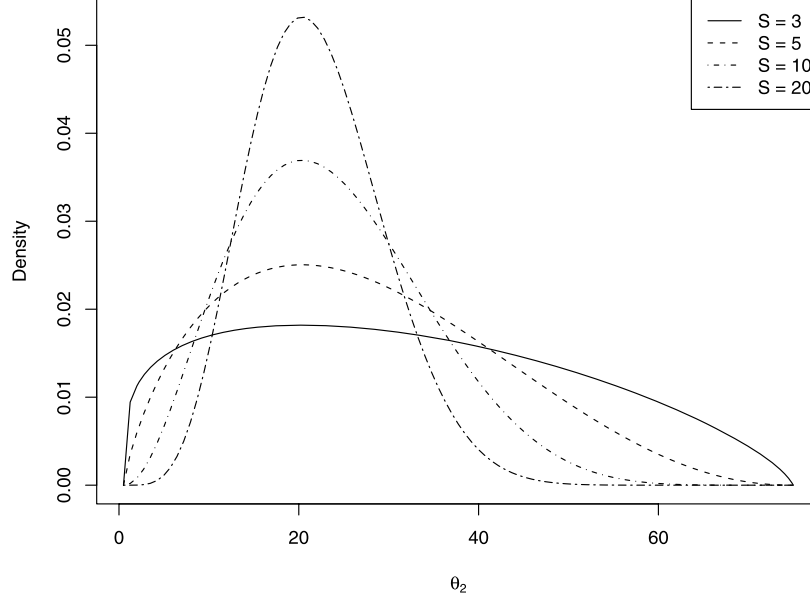
3.2.1. Selection of prior distributions for θ_m . We utilize the factorization in (3.1) to derive a prior distribution for $(\theta_{0m}, \theta_{1m}, \theta_m^0, \sigma^2)$. If θ_m^0 were known, the nonlinear models would reduce to a linear model. It is therefore reasonable to use for a given θ_m^0 the conditionally conjugate normal-inverse gamma (NIG) distribution

$$p(\theta_m^*, \sigma^2 | \theta_m^0) \propto (\sigma^2)^{-(\nu+4)/2} \exp[-\{(\theta_m^* - \mu)' \mathbf{V}^{-1} (\theta_m^* - \mu) + a\} / (2\sigma^2)]$$

for (θ_m^*, σ^2) [see O'Hagan and Forster (2004)], where the parameter θ^* is defined after equation (3.1), $a, \nu > 0$, $\mu \in \mathbb{R}^2$ and $\mathbf{V} \in \mathbb{R}^{2 \times 2}$ denotes a positive definite matrix. The NIG distribution marginally induces a bivariate t -distribution for θ_m^* with ν degrees of freedom, finite mean μ and covariance matrix $a/(\nu-2)\mathbf{V}$, provided that $\nu > 2$. The marginal prior distribution for σ^2 is given by an IG distribution with mode $a/(\nu+2)$, mean $a/(\nu-2)$ and variance $2a^2/\{(\nu-2)^2(\nu-4)\}$. It has a finite mean when $\nu > 2$ and a finite variance when $\nu > 4$.

To set up the NIG distribution for (θ_m^*, σ^2) , one can employ available information about the placebo effect, the maximum treatment effect and the standard deviation. For example, one can choose the marginal bivariate t -distribution for θ_m^* (conditional on θ_m^0) such that the desired mean and covariance are achieved for the placebo effect and the maximum effect of the underlying dose-response model. When the linear parameters θ_m^* cannot be interpreted as placebo and maximum effect, one can use a suitable transformation to achieve the desired moments. Then, one can adjust a and d so that the marginal distribution of σ^2 achieves the desired mode. An attractive choice is to use $\nu = 4$, leading to a prior with infinite variance for σ^2 and a heavy tailed marginal prior for θ_m^* .

For the nonlinear parameters θ_m^0 , we propose selecting suitable bounds for the parameters and then eliciting a bounded prior distribution. This is typically not difficult, as the interpretation of the nonlinear parameters is straightforward, and excessively large parameter values usually correspond to a priori unlikely model shapes. We propose using a scaled beta distribution $B(\alpha, \beta)$ with mode equal to the initial parameter guesses. The curvature of the prior determines the amount of shrinkage that one is willing to employ for the MAP estimates. In the simulations we used the sum $S = \alpha + \beta$ as a measure of curvature with $S > 2$ to ensure unimodality of the beta distribution. Note that already relatively small values of S , such as $S = 10$ or $S = 20$,


 FIG. 2. Beta priors on $[0.5, 75]$ with mode 20 and different S values.

lead to strong shrinkage effects; see Figure 2 for an illustration of the θ_2 parameter in the Emax₁ model, where the initial parameter guess is 20. For dose-response models with more than one nonlinear parameter, we repeat this procedure for all parameters and assume independence among them.

For selecting prior model probabilities, it is convenient to use a uniform distribution across the models unless some models are deemed a priori more plausible than others.

3.2.2. Calculation of posterior probabilities. Let \mathbf{y} denote the data available at an interim analysis and $p(\mathbf{y}|\boldsymbol{\theta}_m, m)$ the likelihood under model m with corresponding prior distribution $p(\boldsymbol{\theta}_m|m)$ and prior model probability $p(m)$. Then the marginal likelihood is given by

$$\begin{aligned} p(m|\mathbf{y}) &\propto p(m) \int p(\mathbf{y}|\boldsymbol{\theta}_m, \sigma^2, m) p(\boldsymbol{\theta}_m, \sigma^2|m) d(\boldsymbol{\theta}_m, \sigma^2) \\ &\propto p(m) \int \int p(\mathbf{y}|\boldsymbol{\theta}_m, \sigma^2, m) p(\boldsymbol{\theta}_m^*, \sigma^2|\boldsymbol{\theta}_m^0, m) d(\boldsymbol{\theta}_m^*, \sigma^2) p(\boldsymbol{\theta}_m^0|m) d\boldsymbol{\theta}_m^0. \end{aligned}$$

The inner integral in the last equation is the product of a likelihood and a conjugate prior distribution. One can hence reduce the integration to

$$(3.5) \quad p(m|\mathbf{y}) \propto p(m) \int p(\mathbf{y}|\boldsymbol{\theta}_m^0, m) p(\boldsymbol{\theta}_m^0|m) d\boldsymbol{\theta}_m^0,$$

where $p(\mathbf{y}|\boldsymbol{\theta}_m^0, m)$ now denotes the integrated likelihood. In our applications, the integral (3.5) is one- or two-dimensional over a bounded region and hence straightforward to calculate numerically. This allows us to calculate the marginal likelihoods efficiently, without resorting to Markov chain Monte Carlo calculations; see Section 3.4 for details. The posterior model probabilities $p(m|\mathbf{y})$ can be obtained by normalizing the marginal likelihoods (multiplied by the prior model probabilities).

We use the maximum $\tilde{\boldsymbol{\theta}}_m^0$ of the marginal posterior $p(\mathbf{y}|\boldsymbol{\theta}_m^0, m)p(\boldsymbol{\theta}_m^0|m)$ as an estimate of $\boldsymbol{\theta}_m^0$. Conditional on this value, we use the maximum $\boldsymbol{\theta}_m^*$ of $p(\boldsymbol{\theta}_m^*|\tilde{\boldsymbol{\theta}}_m^0, \mathbf{y}, m)$ as an estimate for $\boldsymbol{\theta}_m^*$. Therefore, the overall estimate of the parameter $\boldsymbol{\theta}_m$ is given by $\tilde{\boldsymbol{\theta}}_m = (\boldsymbol{\theta}_m^*, \tilde{\boldsymbol{\theta}}_m^0)$. This is a slight variation of the MAP approach described above, but reduces further the computational effort, as it reuses the calculations from the integration to obtain the marginal likelihoods.

3.3. Main algorithm. We now summarize the complete response-adaptive dose-finding design in algorithmic form.

Before trial start:

- (1) Select a starting design using either a balanced allocation across the available doses or an unbalanced allocation based on optimal design considerations.
- (2) Select candidate dose-response models $f_m(d, \boldsymbol{\theta}_m)$.
- (3) Conditional on $\boldsymbol{\theta}_m^0$, calculate a NIG prior distribution for $(\boldsymbol{\theta}_m^*, \sigma^2)$ based on “best guesses” for the placebo effect, the maximum treatment effect and σ^2 (together with suitable variability assumptions for both parameters).
- (4) Choose “best guesses” for the nonlinear parameters $\boldsymbol{\theta}_m^0$ and select the parameter S .
- (5) Choose prior model probabilities $p(m)$ for the different dose-response functions.

At interim analysis:

- (1) Calculate posterior model probabilities

$$(3.6) \quad p(m|\mathbf{y}) \propto p(m) \int p(\mathbf{y}|\boldsymbol{\theta}_m, m)p(\boldsymbol{\theta}_m|m) d\boldsymbol{\theta}_m.$$

Exploiting the conjugacy properties of the NIG distribution, this reduces to one- or two-dimensional integrals; see Section 3.4 for computational details.

- (2) For each dose-response model, estimate $\boldsymbol{\theta}_m$ by using the maximum of $p(\mathbf{y}|\boldsymbol{\theta}_m^0, m)p(\boldsymbol{\theta}_m^0|m)$, where the abscissas calculated in step 1 can be reused. Conditional on this value, use the maximum of $p(\boldsymbol{\theta}_m^*|\tilde{\boldsymbol{\theta}}_m^0, \mathbf{y}, m)$ as an estimate for $\boldsymbol{\theta}_m^*$ to obtain $\tilde{\boldsymbol{\theta}}_m$

(3) Plug the obtained parameter estimates $\tilde{\boldsymbol{\theta}}_m$ into (3.4) and set $\alpha_m = p(m|\mathbf{y})$. Then, minimize with respect to $\mathbf{w} \in \mathcal{S}$, where $\mathcal{S} = \{\mathbf{w} \in \mathbb{S}^k | \mathbf{w} = (\mathbf{n}_{\text{old}} + N_{\text{next}} \mathbf{w}_{\text{next}})(N_{\text{old}} + N_{\text{next}})^{-1}, \mathbf{w}_{\text{next}} \in \mathbb{S}^k\}$. Here, \mathbf{n}_{old} denotes the vector of sample sizes per dose and N_{old} the total sample size until the current interim analysis. Further, N_{next} denotes the sample size and $\mathbf{w}_{\text{next}} \in \mathbb{S}^k$ the design weights for the next cohort of patients. We therefore optimize the design for the next stage taking into account the patient allocation until the current interim analysis; see Section 3.4 for computational details.

(4) Allocate the next cohort of patients according to \mathbf{w}_{next} by applying an appropriate rounding technique, such as described in Pukelsheim (1993), Chapter 12.

Note that the Bayesian approach is used here for design adaptation purposes. The final analysis may or may not be done using a fully Bayesian approach. The development of the Bayesian design methodology above is motivated by the MCP-Mod methodology described in Bretz, Pinheiro and Branson (2005) to address model uncertainty. This method requires prior estimates for the placebo effect, the maximum treatment effect, and σ^2 at the design stage and “best guesses” of the nonlinear parameters $\boldsymbol{\theta}_m^0$ for the analysis. The additional information needed to set up the above adaptive design procedure is hence minimal. Obviously, any other strategy that allows one to use a set of candidate dose-response models might also be used for the final analysis.

3.4. Technical details. In this section we provide further details of the algorithm presented above.

For the calculation of the one- and two-dimensional integrals in (3.5) we used quasi-uniformly distributed point sets based on good lattice points $\mathbf{u}_1, \dots, \mathbf{u}_n$, where $\mathbf{u}_i \in [0, 1]^d$ and d is the dimension of $\boldsymbol{\theta}_m^0$; see Fang and Wang (1994) for details on the construction of such integration grids. Let $\pi(\boldsymbol{\theta}_m^0 | y) = p(\mathbf{y} | \boldsymbol{\theta}_m^0, m) \times p(\boldsymbol{\theta}_m^0 | m)$ denote the integrand from (3.5) and let \mathbf{b}_l and \mathbf{b}_u denote the vector of lower and upper bounds for $\boldsymbol{\theta}_m^0$. One first transforms the good lattice points to obtain $\mathbf{u}_i^* = \mathbf{u}_i(\mathbf{b}_u - \mathbf{b}_l) + \mathbf{b}_l$ for $i = 1, \dots, n$, and then approximates the integral (3.5) by $\prod_{j=1}^d (b_{uj} - b_{lj}) \sum_{i=1}^n \pi(\mathbf{u}_i^* | y) / n$. This approach also allows one to calculate the approximate maximum in the subsequent optimization step by using the grid point \mathbf{u}_i^* corresponding to $\max_i \pi(\mathbf{u}_i^* | y)$. We found that using a grid of size 100 in the one-dimensional case and the good lattice point set of size 1597 in the two-dimensional case [Fang and Wang (1994)] provide sufficiently reliable and computationally efficient results (both for integration and optimization).

The optimization in (3.4) is a constrained optimization problem because the weights \mathbf{w}_{next} lie in the $(k - 1)$ -dimensional probability simplex. A simple but efficient approach to perform the optimization is to use a mapping

$\mathbb{R}^{k-1} \mapsto \mathbb{S}^k$ and then to employ a standard unconstrained optimizer, as described in Atkinson, Donev and Tobias (2007), page 131. To account for the already allocated patients until an interim analysis, one can optimize $\Psi(\frac{n_{\text{old}} + N_{\text{next}} \mathbf{w}_{\text{next}}}{N_{\text{old}} + N_{\text{next}}})$ with respect to $\mathbf{w}_{\text{next}} \in \mathbb{S}^k$. Due to potential multiple optima in the design surface, one cannot be sure whether indeed an optimal design has been found by the optimizer. We thus propose using lower bounds of the resulting relative efficiencies based on the underlying geometry of the optimization problem.

To be precise, suppose that the vector \mathbf{w}^* has been found by the optimizer. The following result gives a lower bound on $r(\mathbf{w}^*) = \frac{\Psi(\mathbf{w}_{\text{opt}})}{\Psi(\mathbf{w}^*)} \in [0, 1]$, where \mathbf{w}_{opt} is the (unknown) true optimal design at the end of the next stage, accounting for the patients allocated until the current interim analysis. A proof of the result is given in the [Appendix](#).

THEOREM 3.1. *A design \mathbf{w} with $c_m(\boldsymbol{\theta}_m) = \nabla b_m(\boldsymbol{\theta}_m) \in \text{Range}(\mathbf{M}_m(\boldsymbol{\theta}_m, \mathbf{w}))$, $m = 1, \dots, M$, minimizes $\Psi(\mathbf{w})$ with respect to \mathbf{w}_{next} , where $\mathbf{w} = \frac{N_{\text{old}} + N_{\text{next}} \mathbf{w}_{\text{next}}}{N_{\text{old}} + N_{\text{next}}}$, if and only if there exist generalized inverses $\mathbf{G}_1, \dots, \mathbf{G}_M$ of $\mathbf{M}_m(\boldsymbol{\theta}_m, \mathbf{w})$, such that the inequality*

$$h(d, \mathbf{w}) = \frac{\sum_{m=1}^M \alpha_m (g_m^T(d, \boldsymbol{\theta}_m) \mathbf{G}_m c_m(\boldsymbol{\theta}_m))^2 / c_m^T(\boldsymbol{\theta}_m) \mathbf{G}_m c_m(\boldsymbol{\theta}_m)}{\sum_{m=1}^M \alpha_m c_m(\boldsymbol{\theta}_m) \mathbf{G}_m^T \mathbf{M}_m(\boldsymbol{\theta}_m, \mathbf{w}_{\text{next}}) \mathbf{G}_m c_m(\boldsymbol{\theta}_m) / c_m^T(\boldsymbol{\theta}_m) \mathbf{G}_m c_m(\boldsymbol{\theta}_m)} \leq 1$$

is satisfied for all $d \in \{d_1, \dots, d_k\}$. Moreover, the efficiency of any design \mathbf{w} can be bounded from below by

$$(3.7) \quad r(\mathbf{w}) \geq \frac{1}{k^*(\mathbf{w}, \gamma)} \geq \frac{1}{h^*(\mathbf{w})},$$

where $\gamma = \frac{N_{\text{old}}}{N_{\text{old}} + N_{\text{next}}}$, $h^*(\mathbf{w}) = \min_{\mathbf{G}_1, \dots, \mathbf{G}_M} \max_{d \in \{d_1, \dots, d_k\}} h(d, \mathbf{w})$,

$$k^*(\mathbf{w}, \gamma) = \left[1 + (1 - \gamma) \min_{\mathbf{G}_1, \dots, \mathbf{G}_M} \max_{\mathbf{v} \in \mathbb{S}^k} \sum_{m=1}^M \alpha_m c_m^T(\boldsymbol{\theta}_m) \mathbf{G}_m^T \times (\mathbf{M}_m(\boldsymbol{\theta}_m, \mathbf{v}) - \mathbf{M}_m(\boldsymbol{\theta}_m, \mathbf{w}_{\text{next}})) \times \mathbf{G}_m c_m(\boldsymbol{\theta}_m) / (c_m^T(\boldsymbol{\theta}_m) \mathbf{G}_m c_m(\boldsymbol{\theta}_m)) \right]^{-1},$$

and the minimum is taken over all generalized inverses $\mathbf{G}_1, \dots, \mathbf{G}_M$ of the matrices $\mathbf{M}_1(\boldsymbol{\theta}_1, \mathbf{w}), \dots, \mathbf{M}_M(\boldsymbol{\theta}_M, \mathbf{w})$.

When the matrices $\mathbf{M}_1(\boldsymbol{\theta}_1, \mathbf{w}), \dots, \mathbf{M}_M(\boldsymbol{\theta}_M, \mathbf{w})$ are invertible, $h^*(\mathbf{w})$ is just the maximum of $h(d, \mathbf{w})$ over the k doses and straightforward to calculate [and so is the lower bound on $r(\mathbf{w})$]. This lower bound is useful in several

respects: we do not need to know the actual optimal design \mathbf{w}_{opt} in order to calculate the lower bound. If the lower bound for our calculated design \mathbf{w}^* is equal to 1, we know that \mathbf{w}^* is the optimal design. Otherwise, we have a conservative estimate on how much percent off one would be when using \mathbf{w} . The bound based on $k^*(\mathbf{w}, \gamma)$ is sharper, however harder to implement.

If one does not use a fully Bayesian approach for the final analysis, one typically has to fit nonlinear regression models to the data. When there are only few doses available, as it is often the case in drug development practice, calculating the ML estimate may be difficult. One way to simplify the problem is by exploiting the fact that θ_0 and θ_1 enter the model function linearly in (3.1). We thus apply the nonlinear optimization only on the nonlinear parameters $\boldsymbol{\theta}_m^0$, similar in spirit to Golub and Pereyra (2003). Using the Frisch–Waugh–Lovell theorem [Baltagi (2008), Chapter 7], we can recalculate the residual sum of squares efficiently, without the need to solve the full least squares problems in each iteration of the nonlinear optimization (this effect becomes even more important when there are additional linear covariates in the model equation, such as gender, baseline values, etc.). In addition, we impose bounds on the nonlinear parameters $\boldsymbol{\theta}_m^0$ to guarantee the existence of the least squares estimate [Seber and Wild (1989), Chapter 12]. As mentioned in Section 3.2.1, such bounds are not a severe restriction in practice and ensure that the optimization problem is well posed.

4. Asthma study revisited. In this section we revisit the asthma case study from Section 2 and address the four open design questions using the proposed methodology from Section 3. To this end, we investigated in an extensive simulation study the operating characteristics for different design options and parameter configurations.

4.1. Design of simulation study. We generated normally distributed observations according to the dose-response models given in Table 1 with $\sigma = 350$ ml. To investigate the robustness of the proposed methods, we also simulated from a linear model (with baseline 100 ml and maximum effect 300 ml) that was not included in the candidate model set. The total sample size was fixed at 300 (constraint imposed by the clinical team). To evaluate the benefit of including additional doses, we compared two design options, one with the four active doses 2.5, 10, 20, 50 (plus placebo) and another one with the seven active doses 0.5, 1, 2.5, 5, 10, 20 and 50 (plus placebo). In addition, we evaluated the benefit of additional interim analyses by varying their number from 0 (= no interim analysis) to 9, where the interim looks were chosen equally spaced in time. In all cases, we assumed a balanced first stage design. The designs from the second stage onward were determined using the observed data according to the algorithm from Section 3.3. When the MED estimate did not exist for certain models at an interim analysis,

they were removed from the model set for the purpose of design calculation and the model probabilities were reweighted accordingly. When the MED estimate did not exist for any model, a balanced allocation was used for the next cohort of patients.

For the final analysis we employed the MCP-Mod procedure from Bretz, Pinheiro and Branson (2005). A potential dose-response signal was assessed using model-based multiple contrast tests based on the candidate model set from Table 1. Subsequently, if there were significant models, the dose-response model with lowest Akaike Information Criterion (AIC) among the significant models was chosen to estimate the MED.

The methodology from Section 3 was applied with uniform prior probabilities for the different models. We further assumed a priori distributions with mean 100 and variance 100,000 for the placebo effect and mean 300 and variance 100,000 for the maximum treatment effect, which were then transformed into the linear parameters for all dose-response models. The mode of the marginal distribution for σ^2 was chosen as 350^2 with $\nu = 4$, resulting in an infinite variance. For the nonlinear parameters we assumed beta distributions (or products thereof) with mode equal to the values specified in Table 1 and $S = 3$. The parameter bounds were chosen to ensure that all reasonable dose-response shapes remained included within the bounds. That is, we chose $\theta_2 \in [0.05, 75]$ for the Emax models, $(\theta_2, \theta_3) \in [0.5, 4] \times [0.5, 4]$ for the beta models and $(\theta_2, \theta_3) \in [0.05, 75] \times [0.5, 25]$ for the the logistic model. For each scenario we used 5000 simulation runs.

4.2. Simulation results. For the chosen standard deviation of $\sigma = 350$ ml, the power of the MCP-Mod procedure to detect a dose-response signal was almost always close to 1. Thus, the MCP-Mod procedure was essentially reduced to choosing the nonlinear model with lowest AIC value under the constraint that only models with significant contrast test statistics were included in the model selection step. Simulations with σ values larger than 350 ml indicated that the power quickly dropped to lower levels (results not shown here), although the estimation results remained qualitatively similar to the ones shown below.

In Figure 3 we display, for each simulation scenario, the mean absolute estimation error for the MED against the number of interim analyses. In all scenarios one observes a benefit from adapting, while most of the improvement is already achieved after 1, 2 or 4 interim analyses. The largest relative improvement (comparing no adaptations vs 9 adaptations) can be observed for the Logistic₁ and the Beta model scenarios, particularly in the case of 7 active doses. The worst relative improvement can be observed for the Logistic₂ scenario, where the overall largest absolute estimation error can be observed. This is not surprising, because even when adapting one cannot achieve a good design for this model, as there are no doses available

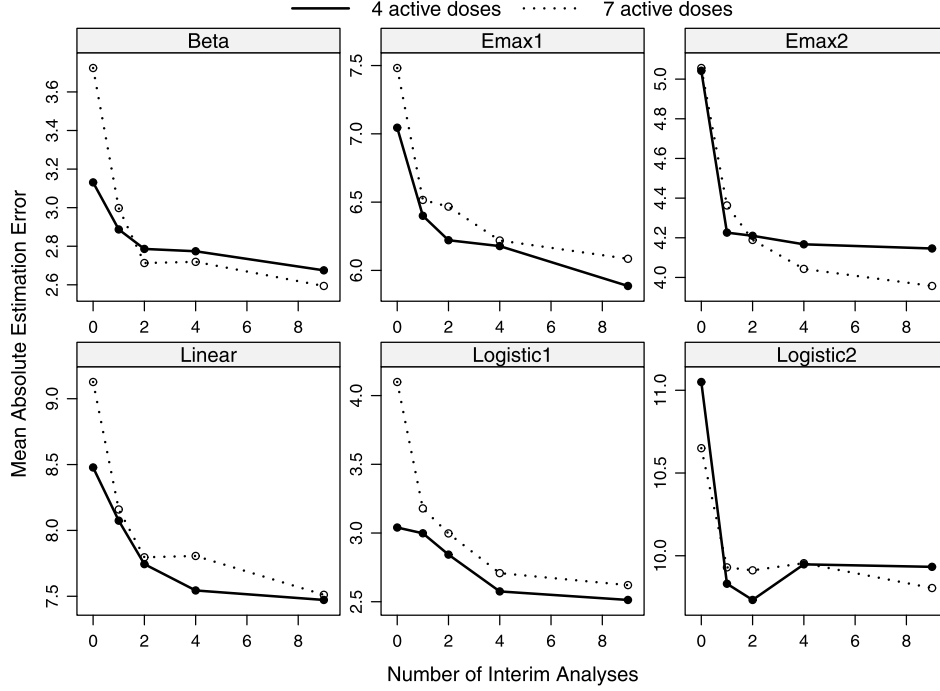


FIG. 3. Mean absolute estimation error for MED estimation.

for administration in the interval $(20, 50)$ containing the MED; see also Figure 1. It is remarkable to see that adaptation also works in the linear model scenario, although the linear model is not included in the candidate model set. It seems that other models in the candidate set are able to capture the shape of a linear model reasonably well.

The comparison between 4 and 7 active doses is not entirely clear. If no interim analyses are performed, it seems that the design with a balanced allocation across the 4 active doses is slightly better than the design with a balanced allocation across all 7 active doses. If one decides to adapt, however, it seems beneficial in some cases to have more doses available, particularly if many interim analyses are performed, while in other cases 4 active doses are sufficient.

To illustrate how adaptation changes the allocation of patients to the different doses, we display in Figure 4 the average patient allocations for the Emax_2 model after 1, 2, 4 and 9 interim analyses and with 7 available active doses. The adaptive design tends to allocate more patients both on placebo and nearby the actual MED. This is intuitively plausible, as the MED estimate depends on the precision of the estimated placebo effect as well as of the estimated function $f(\cdot)$ around the true MED. It also follows from Figure 4 that for a large number of interim analyses the overall

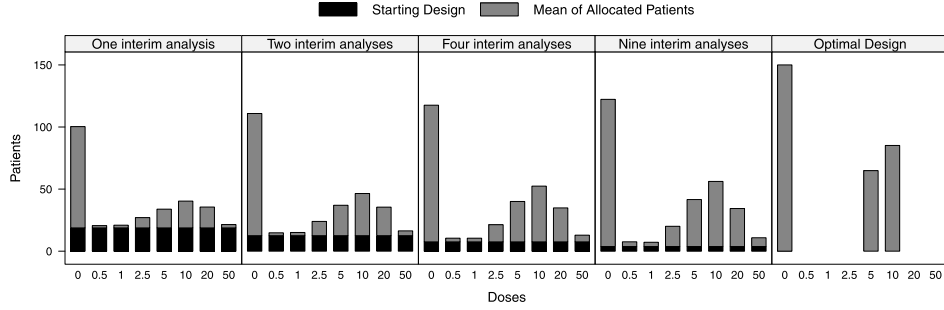


FIG. 4. Average patient allocation after 1, 2, 4 and 9 interim analysis under the E_{\max_2} model. Last panel: locally MED-optimal design for the E_{\max_2} model with true $MED = 7.7$.

allocation is close to the one under a locally optimal design for the E_{\max_2} model, with the variability in the allocations due to the uncertainty both in estimating the correct model and the model parameters at the interim analysis. Similar conclusions also hold for other models than the E_{\max_2} model (not reported here).

We now investigate to which extent the precision gain observed in Figure 3 translates into sample size savings when performing an adaptive design. In other words, how many additional patients are required for a nonadaptive, balanced design to achieve a similar estimation error as with an adaptive design using 300 patients. We again considered the E_{\max_2} model and iterated the total sample size until the mean absolute estimation error was approximately 4 (which is the mean absolute estimation error obtained after 9 interim analyses, as seen in Figure 3). For both design options with 4 and 7 active doses, this was achieved after roughly 500 patients. Thus, using a nonadaptive, balanced design, one would need 200 additional patients to achieve a similar precision in MED estimation as compared to an adaptive design using 300 patients.

The adaptive design benefits observed so far depend on several input parameters, such as the starting design for the first stage. One may argue that starting with a bad design that allocates patients at the “wrong” doses may be improved by adapting at one or more interim looks. On the other hand, starting with a good design may lead to adaptations following random noise at the interim analyses. To illustrate this effect, we report the results for the simulations under the Logistic_1 model (similar results were also obtained for other models and scenarios, but are not reported here). We used four different starting designs. We used $\mathbf{w} = (0.35, 0.03, 0.22, 0.35, 0.05)$ and $\mathbf{w} = (0.35, 0.02, 0.02, 0.02, 0.02, 0.20, 0.30, 0.07)$ as good starting designs with 4 and 7 active doses, respectively. These designs work well because they allocate patients on placebo and around the MED, while keeping some mass on the remaining doses. In addition, we used $\mathbf{w} = (0.1, 0.3, 0.05, 0.05, 0.5)$ and $\mathbf{w} = (0.1, 0.2, 0.22, 0.02, 0.02, 0.02, 0.02, 0.4)'$ as bad starting designs, as

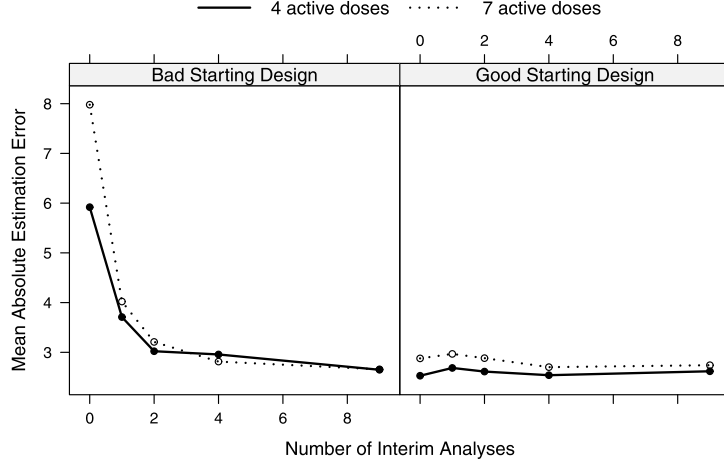


FIG. 5. Mean absolute estimation error for MED estimation, under the the Logistic_1 model for good and bad starting designs.

they have relatively few patients on placebo and around the MED. It follows from Figure 5 that substantial improvements are possible when using bad starting designs. On the other hand, for good starting designs no benefit is achieved by adapting and the performance may even deteriorate, because the possibility of adapting may lead one to deviate from the already good starting design. In practice, one does not know whether an employed design is good or bad, but one should keep in mind the possibility that adaptive designs will not always improve upon the initial design.

To further investigate the robustness of the proposed methods, we repeated the simulation study from Figure 3 by increasing the standard deviation to 450 ml and 700 ml. The overall results remain similar, but with increased absolute estimation errors. However, the relative benefit of 9 interim analyses vs. no adaptation decreases slightly. Due to the larger noise, one obtains less reliable information at an interim analysis and one may end up with a worse design for the next stage. We also investigated the effect of prior misspecification. For this purpose we misspecified the prior means or prior modes by adding or subtracting 20% of the true value, but leaving the variability (variance of baseline and maximum effect and the value S for the beta distribution) unchanged as in the original simulations. The results are largely identical to those presented in Figure 3, indicating that the proposed methods are robust under moderate prior misspecifications.

4.3. Conclusions for asthma study. Many more simulations than presented above were conducted at the planning stage of the asthma study to address the four questions stated in Section 2. Regarding question (A), it was felt that the potential benefits of conducting an adaptive design (more

precise MED estimation) outweighed the additional logistical requirements, especially in view of the perceived sample size gain of 100–200 patients when compared to a fixed-sample study designed to achieve a similar precision. For question (B) it was decided to have one interim analysis: based on Figure 3 and other simulation results, the potential further reduction of the mean absolute estimation error with two or more interim analyses was perceived as too small to justify the additional logistical complexity.

For similar reasons, it was decided against having all seven active doses from the beginning on question (C). Instead, 150 patients ought to be allocated equally across the four active doses 2.5, 10, 20, 50 (plus placebo) in the first stage. Once the interim results are available and analyzed with the methods from Section 3, however, patients could be allocated to all seven active doses (or a subset thereof) in the second stage. For practical reasons, the clinical team decided to incorporate constraints on the minimum number of patients allocated per dose in the second stage: if the algorithm would allocate less than 5% of the patients on a certain dose, that dose would be dropped altogether and the corresponding patients reallocated to the remaining doses.

5. Discussion. Motivated by a dose finding study in moderate persistent asthma, we described a response-adaptive approach that addresses common challenges encountered in dose-finding studies: multiple objectives, model uncertainty, and large variability. When planning an adaptive dose-finding design it is important to realize that it may not always be better than a nonadaptive design. It is necessary to employ a factored view, as many parameters may impact the performance of a study design. Often, an unbalanced fixed-sample design derived from optimal design theory might already provide benefit over a balanced fixed-sample design and adaptation may not bring further advantages, particularly if the variability is large (which is common in practice). Thus, adaptive designs are promising in situations where the initial design is not good and/or interim parameter estimates have low variability. In practice, one never knows how good the initial design will be, before trial start, and adaptive designs may guard against bad initial designs. However, the benefits of adaptive dose-finding designs have to be balanced against the increased logistical requirements to implement processes for repeated data collection, cleaning and analyses, to maintain trial integrity and validity, and to overcome potential challenges in drug manufacturing and supply.

In this paper we focused on designs based on the compound optimality criterion (3.4) to address model uncertainty and to minimize the variance of $\widehat{\text{MED}}$. The criterion depends on the parameters of the different dose-response models as well as on the model probabilities and we used

a Bayesian approach to continuously update parameter values and model probabilities based on the information accrued in the trial. The approach was implemented based on optimization and numerical quadrature, so that computationally intensive Markov chain Monte Carlo techniques could be avoided. Computational efficiency is of extreme importance, as the frequentist operating characteristics of any adaptive design methodology needs to be evaluated in extensive simulations under multiple scenarios.

The proposed method can be extended immediately if alternative optimality criteria [such as ED_p - or D-optimal designs, see Dette et al. (2010)] or mixtures thereof are of interest. Alternatively, optimal discrimination designs could be applied that allow one to differentiate among several candidate nonlinear regression models [Atkinson and Fedorov (1975), Dette and Titoff (2009)]. It would be interesting to address multiple objectives by considering different optimality criteria at different stages, such as using a model discrimination design in earlier stages, and MED-optimal design in later stages. This will be investigated in future research, but see Dragalin et al. (2010) for initial results.

The R functions used for the simulations are available with the `DoseFinding` R package [see Bornkamp, Pinheiro and Bretz (2010)].

APPENDIX: PROOF OF THEOREM 3.1

Obviously the first part of the theorem follows from the lower bound (3.7) on the efficiency. For a proof of (3.7) let $\gamma = N_{\text{old}}/(N_{\text{old}} + N_{\text{next}}) \in (0, 1)$ and note that the total information of the experiment in the m th model is given by

$$(A.1) \quad \mathbf{M}_m(\boldsymbol{\theta}, \mathbf{w}_{\text{old}}, \mathbf{w}_{\text{next}}) = \gamma \mathbf{M}_m(\boldsymbol{\theta}_m, \mathbf{w}_{\text{old}}) + (1 - \gamma) \mathbf{M}_m(\boldsymbol{\theta}_m, \mathbf{w}_{\text{next}}),$$

where we collect in the vector $\boldsymbol{\theta} = (\boldsymbol{\theta}_1, \dots, \boldsymbol{\theta}_M)$ the parameters of the different models. Define a block diagonal matrix by

$$(A.2) \quad \begin{aligned} &\overline{\mathbf{M}}(\boldsymbol{\theta}_m, \mathbf{w}_{\text{old}}, \mathbf{w}_{\text{next}}) \\ &= \text{diag}(\mathbf{M}_1(\boldsymbol{\theta}_m, \mathbf{w}_{\text{old}}, \mathbf{w}_{\text{next}}), \dots, \mathbf{M}_M(\boldsymbol{\theta}_m, \mathbf{w}_{\text{old}}, \mathbf{w}_{\text{next}})) \end{aligned}$$

(all other entries in this matrix are 0) and, similarly,

$$K = \text{diag}(c_1(\boldsymbol{\theta}_1), \dots, c_M(\boldsymbol{\theta}_M)),$$

where the vector $c_m(\boldsymbol{\theta})$ is given by $c_m(\boldsymbol{\theta}) = \nabla b_m(\boldsymbol{\theta})$, $m = 1, \dots, M$. For a design \mathbf{w}_{next} , such that $c_m(\boldsymbol{\theta}) \in \text{Range}(\mathbf{M}_m(\boldsymbol{\theta}_m, \mathbf{w}_{\text{old}}, \mathbf{w}_{\text{next}}))$ ($m = 1, \dots, M$), we consider the information matrix

$$\begin{aligned} C_K(\overline{\mathbf{M}}(\boldsymbol{\theta}, \mathbf{w}_{\text{old}}, \mathbf{w}_{\text{next}})) &= (K^T \overline{\mathbf{M}}^{-1}(\boldsymbol{\theta}, \mathbf{w}_{\text{old}}, \mathbf{w}_{\text{next}}) K)^{-1} \\ &= \text{diag}((c_1^T(\boldsymbol{\theta}_1) \mathbf{M}_1^{-1}(\boldsymbol{\theta}_1, \mathbf{w}_{\text{old}}, \mathbf{w}_{\text{next}}) c_1(\boldsymbol{\theta}_1))^{-1}, \dots, \\ &\quad (c_M^T(\boldsymbol{\theta}_M) \mathbf{M}_M^{-1}(\boldsymbol{\theta}_M, \mathbf{w}_{\text{old}}, \mathbf{w}_{\text{next}}) c_M(\boldsymbol{\theta}_M))^{-1}). \end{aligned}$$

Note that the optimal design maximizes

$$\begin{aligned}\Psi^{-1}(\mathbf{w}_{\text{next}}) &= \frac{N_{\text{old}} + N_{\text{next}}}{\sigma^2} \cdot \Phi_{\alpha}(C_K(\overline{\mathbf{M}}(\boldsymbol{\theta}, \mathbf{w}_{\text{old}}, \mathbf{w}_{\text{next}}))) \\ &= \frac{N_{\text{old}} + N_{\text{next}}}{\sigma^2} \prod_{m=1}^M (c_m^T(\boldsymbol{\theta}) \mathbf{M}_M^{-}(\boldsymbol{\theta}_m, \mathbf{w}_{\text{old}}, \mathbf{w}_{\text{next}}) c_m(\boldsymbol{\theta}))^{-\alpha_m},\end{aligned}$$

where the last identity defines the criterion Φ_{α} and we have used the notation $\Phi_{\alpha}(\text{diag}(\lambda_1, \dots, \lambda_M)) = \prod_{m=1}^M \lambda_m^{\alpha_m}$. Now according to Theorem 1 in Dette (1996), a lower bound for the efficiency of the design \mathbf{w}_{next}

$$r(\mathbf{w}) = \frac{\psi^{-1}(\mathbf{w})}{\psi^{-1}(\mathbf{w}_{\text{opt}})} = \frac{\Phi_{\alpha}(C_K(\overline{\mathbf{M}}(\boldsymbol{\theta}, \mathbf{w}_{\text{old}}, \mathbf{w}_{\text{next}})))}{\max_{\mathbf{v} \in \mathbb{S}^k} \Phi_{\alpha}(C_K(\overline{\mathbf{M}}(\boldsymbol{\theta}, \mathbf{w}_{\text{old}}, \mathbf{v})))}$$

is obtained as

$$\begin{aligned}e &= \left[\min_{\mathbf{G}} \max_{A \in \mathcal{A}} \text{tr}\{\mathbf{G} K C_K(\overline{\mathbf{M}}(\boldsymbol{\theta}, \mathbf{w}_{\text{old}}, \mathbf{w}_{\text{next}})) \right. \\ &\quad \left. \times \mathbf{E} C_K(\overline{\mathbf{M}}(\boldsymbol{\theta}, \mathbf{w}_{\text{old}}, \mathbf{w}_{\text{next}})) K^T \mathbf{G}^T A\} \right]^{-1},\end{aligned}$$

where the minimum is taken over the set of all generalized inverses of the matrix $\overline{\mathbf{M}}(\boldsymbol{\theta}, \mathbf{w}_{\text{old}}, \mathbf{w}_{\text{next}})$ and the set \mathcal{A} is defined by

$$\mathcal{A} = \{\overline{\mathbf{M}}(\boldsymbol{\theta}, \mathbf{w}_{\text{old}}, \mathbf{v}) \mid \mathbf{v} \in \mathbb{S}^k\}$$

and the matrix \mathbf{E} is given by

$$\mathbf{E} = \text{diag}(\alpha_1 c_1^T(\boldsymbol{\theta}) \mathbf{M}_1^{-}(\boldsymbol{\theta}_1, \mathbf{w}_{\text{old}}, \mathbf{v}) c_1(\boldsymbol{\theta}), \dots, \alpha_M c_M^T(\boldsymbol{\theta}) \mathbf{M}_M^{-}(\boldsymbol{\theta}_M, \mathbf{w}_{\text{old}}, \mathbf{v}) c_M(\boldsymbol{\theta})).$$

Therefore, observing the identity

$$\mathbf{M}_m(\boldsymbol{\theta}_m, \mathbf{w}_{\text{old}}, \mathbf{v}) = \mathbf{M}_m(\boldsymbol{\theta}_m, \mathbf{w}_{\text{old}}, \mathbf{w}_{\text{next}}) + (1 - \gamma)(\mathbf{M}(\boldsymbol{\theta}_m, \mathbf{v}) - \mathbf{M}(\boldsymbol{\theta}_m, \mathbf{w}_{\text{next}})),$$

we obtain

$$\begin{aligned}e &= \left[1 + (1 - \gamma) \min_{\mathbf{G}_m} \max_{v \in \mathbb{S}^k} \sum_{m=1}^M \alpha_m c_m^T(\boldsymbol{\theta}_m) \mathbf{G}_m^T (\mathbf{M}_m(\boldsymbol{\theta}_m, \mathbf{v}) - \mathbf{M}_m(\boldsymbol{\theta}_m, \mathbf{w}_{\text{next}})) \right. \\ &\quad \left. \times \mathbf{G}_m c_m(\boldsymbol{\theta}_m) / (c_m^T(\boldsymbol{\theta}_m) \mathbf{G}_m c_m(\boldsymbol{\theta}_m)) \right]^{-1} \\ &\geq \left[\min_{\mathbf{G}_1, \dots, \mathbf{G}_M} \max_{d \in \{d_1, \dots, d_k\}} \sum_{m=1}^M \alpha_m \frac{(g_m^T(d, \boldsymbol{\theta}_m) \mathbf{G}_m c_m(\boldsymbol{\theta}_m))^2}{c_m^T(\boldsymbol{\theta}_m) \mathbf{G}_m c_m(\boldsymbol{\theta}_m)} \right. \\ &\quad \left. / \sum_{m=1}^M \alpha_m \frac{c_m^T(\boldsymbol{\theta}_m) \mathbf{G}_m^T \mathbf{M}_m(\boldsymbol{\theta}_m, \mathbf{w}_{\text{next}}) \mathbf{G}_m c_m(\boldsymbol{\theta}_m)}{c_m^T(\boldsymbol{\theta}_m) \mathbf{G}_m c_m(\boldsymbol{\theta}_m)} \right]^{-1},\end{aligned}$$

where we have used the inequality

$$[1 + (1 - \gamma)(A - B)]^{-1} \geq \left[\frac{A}{B} \right]^{-1}$$

for $A \geq B \geq 0$, $(1 - \gamma)B \leq 1$ and standard arguments in design theory.

Acknowledgments. The authors would like to thank Martina Stein, who typed parts of this manuscript with considerable technical expertise.

REFERENCES

- ATKINSON, A. C., DONEV, A. N. and TOBIAS, R. D. (2007). *Optimum Experimental Designs, with SAS. Oxford Statistical Science Series* **34**. Oxford Univ. Press, Oxford. [MR2323647](#)
- ATKINSON, A. C. and FEDOROV, V. V. (1975). Optimal design: Experiments for discriminating between several models. *Biometrika* **62** 289–303. [MR0381163](#)
- BALTAGI, B. H. (2008). *Econometrics*, 4th ed. Springer, Berlin.
- BERGER, J. O. and PERICCHI, L. R. (1996). The intrinsic Bayes factor for model selection and prediction. *J. Amer. Statist. Assoc.* **91** 109–122. [MR1394065](#)
- BIEDERMANN, S., DETTE, H. and PEPELYSHEV, A. (2006). Some robust design strategies for percentile estimation in binary response models. *Canad. J. Statist.* **34** 603–622. [MR2347048](#)
- BORNKAMP, B. (2006). Comparison of model-based and model-free approaches for the analysis of dose-response studies. Diploma thesis, Fakultät Statistik, Technische Univ. Dortmund. Available at <http://www.statistik.uni-dortmund.de/~bornkamp/diplom.pdf>.
- BORNKAMP, B., PINHEIRO, J. and BRETZ, F. (2010). DoseFinding: Planning and analyzing dose finding experiments. R package version 0.4-1.
- BRETZ, F., PINHEIRO, J. C. and BRANSON, M. (2005). Combining multiple comparisons and modeling techniques in dose-response studies. *Biometrics* **61** 738–748. [MR2196162](#)
- BRETZ, F., HSU, J., PINHEIRO, J. and LIU, Y. (2008). Dose finding—a challenge in statistics. *Biom. J.* **50** 480–504. [MR2526514](#)
- CHALONER, K. and VERDINELLI, I. (1995). Bayesian experimental design: A review. *Statist. Sci.* **10** 273–304. [MR1390519](#)
- CHATFIELD, C. (1995). Model uncertainty, data mining and statistical inference. *J. Roy. Statist. Soc. Ser. A* **158** 419–466.
- COOK, R. D. and WONG, W. K. (1994). On the equivalence of constrained and compound optimal designs. *J. Amer. Statist. Assoc.* **89** 687–692. [MR1294092](#)
- DETTE, H. (1996). Lower bounds for efficiencies with applications. In *Research Developments in Probability and Statistics* 111–124. VSP, Utrecht. [MR1462412](#)
- DETTE, H. and TITOFF, S. (2009). Optimal discrimination designs. *Ann. Statist.* **37** 2056–2082. [MR2533479](#)
- DETTE, H., BRETZ, F., PEPELYSHEV, A. and PINHEIRO, J. (2008). Optimal designs for dose-finding studies. *J. Amer. Statist. Assoc.* **103** 1225–1237. [MR2462895](#)
- DETTE, H., KISS, C., BEVANDA, M. and BRETZ, F. (2010). Optimal designs for the emax, log-linear and exponential models. *Biometrika* **97** 513–518. [MR2650755](#)
- DRAGALIN, V., HSUAN, F. and PADMANABHAN, S. K. (2007). Adaptive designs for dose-finding studies based on sigmoid E_{\max} model. *J. Biopharm. Statist.* **17** 1051–1070. [MR2414561](#)

- DRAGALIN, V., BORNKAMP, B., BRETZ, F., MILLER, F., PADMANABHAN, S. K., PATEL, N., PEREVOZSKAYA, I., PINHEIRO, J. and SMITH, J. R. (2010). A simulation study to compare new adaptive dose-ranging designs. *Statistics in Biopharmaceutical Research* **2** 487–512.
- DRAPER, D. (1995). Assessment and propagation of model uncertainty. *J. Roy. Statist. Soc. Ser. B* **57** 45–97. With discussion and a reply by the author. [MR1325378](#)
- FANG, K. T. and WANG, Y. (1994). *Number-theoretic Methods in Statistics. Monographs on Statistics and Applied Probability* **51**. Chapman & Hall, London. [MR1284470](#)
- FEDOROV, V. and LEONOV, S. (2001). Optimal design of dose response experiments: A model-oriented approach. *Drug Information Journal* **35** 1373–1383.
- GOLUB, G. and PEREYRA, V. (2003). Separable nonlinear least squares: The variable projection method and its applications. *Inverse Problems* **19** R1–R26. [MR1991786](#)
- HJORTH, J. S. U. (1994). *Computer Intensive Statistical Methods: Validation, Model Selection, and Bootstrap*. Chapman & Hall, London. [MR1270448](#)
- KASS, R. and RAFTERY, A. (1995). Bayes factors. *J. Amer. Statist. Assoc.* **90** 773–795.
- KING, J. and WONG, W. K. (2004). Optimal designs for the power logistic model. *J. Statist. Comput. Simul.* **74** 779–791. [MR2083103](#)
- LÄUTER, E. (1974). Experimental design in a class of models. *Math. Operationsforsch. Statist.* **5** 379–398. [MR0440812](#)
- MILLER, F., GUILBAUD, O. and DETTE, H. (2007). Optimal designs for estimating the interesting part of a dose-effect curve. *J. Biopharm. Statist.* **17** 1097–1115. [MR2414564](#)
- MÜLLER, P., BERRY, D. A., GRIEVE, A. P. and KRAMS, M. (2006). A Bayesian decision-theoretic dose-finding trial. *Decision Analysis* **3** 197–207.
- O’HAGAN, A. (1995). Fractional Bayes factors for model comparison. *J. Roy. Statist. Soc. Ser. B* **57** 99–138. With discussion and a reply by the author. [MR1325379](#)
- O’HAGAN, A. and FORSTER, J. (2004). *Kendall’s Advanced Theory of Statistics, Volume 2B: Bayesian Inference*, 2nd ed. Arnold, London.
- PINHEIRO, J., BORNKAMP, B. and BRETZ, F. (2006). Design and analysis of dose-finding studies combining multiple comparisons and modeling procedures. *J. Biopharm. Statist.* **16** 639–656. [MR2252312](#)
- PUKELSHEIM, F. (1993). *Optimal Design of Experiments*. Wiley, New York. [MR1211416](#)
- RUBERG, S. J. (1995). Dose response studies. I. Some design considerations. *J. Biopharm. Statist.* **5** 1–14.
- SEBER, G. A. F. and WILD, C. J. (1989). *Nonlinear Regression*. Wiley, New York. [MR0986070](#)
- VAN DER VAART, A. W. (1998). *Asymptotic Statistics. Cambridge Series in Statistical and Probabilistic Mathematics* **3**. Cambridge Univ. Press, Cambridge. [MR1652247](#)
- WATHEN, J. K. and THALL, P. F. (2008). Bayesian adaptive model selection for optimizing group sequential clinical trials. *Statist. Med.* **27** 5586–5604. [MR2573771](#)
- WU, C. F. J. (1988). Optimal design for percentile estimation of a quantal response curve. In *Optimal Design and Analysis of Experiments* 213–222. Elsevier, Amsterdam.
- ZHOU, X., JOSEPH, L., WOLFSON, D. B. and BÉLISLE, P. (2003). A Bayesian A-optimal and model robust design criterion. *Biometrics* **59** 1082–1088. [MR2025133](#)
- ZHU, W. and WONG, W. K. (2000). Multiple-objective designs in a dose-response experiment. *J. Biopharm. Statist.* **10** 1–14.
- ZHU, W. and WONG, W. K. (2001). Bayesian optimal designs for estimating a set of symmetric quantiles. *Stat. Med.* **20** 123–137.

B. BORNKAMP
TECHNISCHE UNIVERSITÄT DORTMUND
FAKULTÄT STATISTIK
44780 DORTMUND
GERMANY
E-MAIL: bornkamp@statistik.tu-dortmund.de

H. DETTE
RUHR-UNIVERSITÄT BOCHUM
FAKULTÄT FÜR MATHEMATIK
44780 BOCHUM
GERMANY
E-MAIL: holger.dette@rub.de

F. BRETZ
NOVARTIS PHARMA AG
STATISTICAL METHODOLOGY
4002 BASEL
SWITZERLAND
E-MAIL: frank.bretz@novartis.com

J. PINHEIRO
JOHNSON & JOHNSON
RARITAN, NEW JERSEY 08869
USA
E-MAIL: jpinhei1@its.jnj.com